

On the Capacity Bounds of Undirected Networks

Ali Al-Bashabsheh and Abbas Yongacoglu
 School of Information Technology and Engineering
 University of Ottawa, Ottawa, Canada
 {aalba059, yongacog}@site.uottawa.ca

Abstract—In this work we improve on the bounds presented in [1] for network coding gain in the undirected case. A tightened bound for the undirected multicast problem with three terminals is derived. An interesting result shows that with fractional routing, routing throughput can achieve at least 75% of the coding throughput. A tighter bound for the general multicast problem with any number of terminals shows that coding gain is strictly less than 2. Our derived bound depends on the number of terminals in the multicast network and approaches 2 for arbitrarily large number of terminals.

I. INTRODUCTION

Network coding gain, \mathfrak{G} , is defined as the improvement in network throughput due to coding compared to routing throughput. It is well known that network coding increases the space over which throughput is maximized and thus provides a capacity at least as high as the routing one. In this work, capacity of a network refers to the maximum achievable throughput under a certain coding scheme (routing can be considered as a coding scheme with identity mapping). For a set of sinks in a directed multicast network, it was shown in [2] that if the network can achieve a certain throughput to each receiver individually, then it can achieve the same throughput to all the sinks simultaneously by allowing coding at intermediate nodes and thus, achieve a throughput gain. This has ignited an area of research trying to answer many questions; one of them is how much gain is possible? There has been some instances in the literature where the network coding gain can be unbounded for directed networks [3] [4]. Relating coding gain to the integrality gap of linear programming formulation for minimum weight Steiner tree [5], further examples of directed networks with arbitrarily high coding gain can be obtained [6] [7]. For undirected networks, it was shown in [1] that network coding gain is bounded by 2 when half integer routing is possible. Such a dramatic difference between directed and undirected multicast networks results because directed networks can be oriented in a way such that routing can offer little compared to coding. In this work we further investigate achievable routing throughputs and derive tighter bounds on the possible coding gain.

This work is organized as follows. Section II introduces some necessary definitions. In Section III we use a Steiner tree packing argue to present bounds on the achievable routing throughput and derive upper bounds on coding gain.

II. DEFINITIONS

An undirected network G on V nodes and E links can be modeled as an undirected graph $G(V, E)$ that might contain

multiple edges, i.e. a multigraph. At certain places we might refer to the sets of vertices and edges of a graph H (especially if they are not explicitly specified) as $V(H)$ and $E(H)$, respectively. A multicast network with a source node $s \in V$ and a set of sinks $\mathcal{R} = \{R_1, \dots, R_{|\mathcal{R}|}\} \subseteq V - s$ is a network where s broadcasts data to all the receivers in \mathcal{R} . At certain places in this work we might not distinguish between a source or a sink node and simply refer to such a node as a *terminal*. If $|\mathcal{R}| = V - 1$ then we call G a *broadcast* network.

We assume that edges capacities are defined over the same base as the source symbols. In other words, an edge with capacity C can carry at most C symbols. Integer routing throughput refers to the routing throughput achieved by routing scalar symbols through edges of the network. Considering multiple time units in a network, fractional routing throughput denotes the throughput achieved by upscaling edge capacities by n divided by n . Such a routing scheme may also be referred to as vector routing and it was shown in [8] that doubling edges capacities can result in a throughput more than twice as good. If we let h denote the number of symbols a multicast network can deliver from the source to all sinks, we can define coding and fractional and integer routing capacities as:

$$\begin{aligned}\gamma &= \sup \{h/n \in \mathbb{Q}^+ : h/n \text{ is an achievable coding rate}\} \\ \pi_f &= \sup \{h/n \in \mathbb{Q}^+ : h/n \text{ is an achievable routing rate}\} \\ \pi_i &= \sup \{h \in \mathbb{Z}^+ : h \text{ is an achievable integer routing rate}\}\end{aligned}$$

where \mathbb{Z}^+ and \mathbb{Q}^+ are the sets of positive integers and positive rational numbers, respectively. A fractional routing scheme with $n = 2$ will be referred to as half integer routing and the corresponding capacity will be denoted by $\pi_{\frac{1}{2}}$. In the following, if a statement is true for both integer and fractional routing, we may drop the subscripts and simply write π . The gain in throughput due to coding compared to integer, half-integer and fractional routing is defined as $\mathfrak{G}_i = \frac{\gamma}{\pi_i}$, $\mathfrak{G}_{\frac{1}{2}} = \frac{\gamma}{\pi_{\frac{1}{2}}}$ and $\mathfrak{G}_f = \frac{\gamma}{\pi_f}$, respectively.

For a graph $G = (V, E)$, an induced subgraph T is called a *spanning tree* if and only if T is a tree and $V(T) = V(G)$. For $A \subseteq V$, a subtree T of G is called an *A-Steiner tree* (A-Spanning tree) if and only if $A \subseteq V(T)$. Steiner tree packing refers to finding the maximum number of edge disjoint A-Steiner trees in G .

Edge connectivity, $\lambda(G)$, of a connected graph G is the size of the smallest set $F \subseteq E(G)$ such that $G - F$ is disconnected ($G - F$ denotes the graph induced from G by deleting all edges $e \in F$). A graph G is l -edge connected for any integer $l \leq$

$\lambda(G)$. For $A \subseteq V(G)$, $\lambda_G(A)$ denotes the edge connectivity of A in G . The minimum size of a cut between a pair of vertices $u, v \in V(G)$ is the edge connectivity, $\lambda_G(u, v)$, of u and v in G which by Menger's theorem equals the maximum number of edge disjoint paths between u and v . A *cut-edge* (bridge) is an edge $e \in E(G)$ such that $G - e$ is disconnected (clearly, $\lambda(G)$ must be 1 for such an edge to exist). The degree of $u \in V$ is the number of incident edges to u and is denoted as $d(u)$. If the underlying graph G is clear from the context, we might drop the subscript G from the previous quantities.

Let $e = rx$ and $f = xt$ be two edges in $G(V, E)$. *Splitting off* the pair of edges e and f refers to deleting e and f and introducing a new edge $g = rt$. We refer to the pair e, f and g as the splitted and the splitting edges, respectively. We also denote the resulting graph after splitting off edges e, f as G^{ef} . A pair of edges e, f incident with x is *admissible* if $\lambda_{G^{ef}}(u, v) = \lambda_G(u, v)$ for every $u \neq v \in V - x$. An incident edge to x is admissible if it belongs to an admissible pair, otherwise it is non-admissible. A complete splitting at $x \in V$ (when $d(x)$ is even) refers to: 1) repeatedly splitting off pairs of incident edges to x until x is isolated, and 2) deleting x . We denote the graph resulting from a complete splitting at x as G^x and call such a splitting a *suitable* one if every splitting in step 1) is admissible, i.e. all edges incident to x belong to *disjoint* admissible pairs.

III. BOUNDS ON ACHIEVABLE RATES FOR UNDIRECTED NETWORKS

For network G modeled as an undirected graph $G(V, E)$ and a set of terminals $A \subseteq V$, using a max-flow min-cut argument, the network coding capacity is trivially upper bounded by $\lambda(A)$. Note that unlike directed networks (where the min-cut throughput is always achievable via coding over a sufficiently large field [2] [9]) in the undirected case a throughput of $\lambda(A)$ might not be achievable (see example 2). It is also clear that $\pi \leq \gamma$ since routing is a special type of coding where only repetition and forwarding are allowed at relay nodes. If $\lambda(A) = 1$, then it is clear that $\gamma = \pi = 1$ and hence we can always assume $\lambda(A) \geq 2$. On the other hand, if $\lambda(G) = 1$ and $\lambda(A) \geq 2$ then $V(G)$ can be partitioned into two disjoint subsets $V(G_1)$ and $V(G_2)$ with $A \subseteq V(G_1)$ and G_1 and G_2 connected via a cut-edge. In this case we can always delete G_2 without affecting the throughput. In conclusion, we can always assume that G is 2-edge connected, i.e. G contains no cut-edges.

A. Multicast Networks with Three Terminals

Let $G = (V, E)$ be a graph representing a multicast network with a set of terminals $A = \{s, R_1, R_2\} \subseteq V$, the following theorem provides a lower bound on the routing capacity,

Theorem 1: *For an undirected multicast network with three terminals, the integer, half integer and fractional routing capacities are bounded as*

$$\pi_i \geq \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor$$

$$\begin{aligned} \pi_{\frac{1}{2}} &\geq \frac{1}{2} \left\lfloor \frac{12\lambda(A) - 3}{8} \right\rfloor \\ \pi_f &\geq \frac{3\lambda(A)}{4} - \epsilon \end{aligned}$$

where ϵ is arbitrarily small for arbitrarily large n and $\lambda(A)$ is the connectivity of A in G .

Proof: The proof makes use of the following theorem

Theorem 2 (Kriesell, [10]). *For any integer $k \geq 1$, let $A = \{v_1, v_2, v_3\}$ be $\lfloor \frac{8k+3}{6} \rfloor$ -edge connected in G , then there exists a system of k edge disjoint A -Steiner trees in G .*

Let $A = \{v_1, v_2, v_3\}$ be the set of terminals and let A be $\lambda(A)$ -edge connected in G . Noting that any positive integer, a , can be written as

$$\begin{aligned} a &= dq + b; \text{ for some choice } q \in \{0\} \cup \mathbb{Z}^+, \\ d &\in \{1, 2, \dots, a\} \text{ and } b \in \{0, 1, 2, \dots, d-1\} \end{aligned} \quad (1)$$

Consider,

$$\left\lfloor \frac{8k+3}{6} \right\rfloor = \left\lfloor \frac{(6+2)k+3}{6} \right\rfloor = k + \left\lfloor \frac{2k+3}{6} \right\rfloor, \quad \forall k \in \mathbb{Z}^+$$

From (1) with $d = 3$, we can write $k = 3q + m$, $m \in \{0, 1, 2\}$. Thus,

$$\begin{aligned} \left\lfloor \frac{8k+3}{6} \right\rfloor &= 3q + m + \left\lfloor \frac{6q+2m+3}{6} \right\rfloor \\ &= 4q + m + \left\lfloor \frac{2m+3}{6} \right\rfloor, \quad m \in \{0, 1, 2\} \\ &= \begin{cases} 4q, & m = 0 \\ 4q + 1, & m = 1 \\ 4q + 3, & m = 2 \end{cases} \end{aligned}$$

Thus, if $\lambda(A)$ is any positive integer such that $\lambda(A) = 4q + j$, $j \in \{0, 1, 3\}$, we can write $\lambda(A) = \lfloor \frac{8k+3}{6} \rfloor$ for some positive integer k . The only class of positive integers left is the one such that $\lambda(A) = 4q + 2$. For this choice of integers, we can write $\lambda(A) = 4q + 1 + 1 = \lfloor \frac{8k+3}{6} \rfloor + 1$ for some positive integer k . Hence, for any positive integer, $\lambda(A)$, there exist a positive integer k for which $\lambda(A) = \lfloor \frac{8k+3}{6} \rfloor + \delta$, where k is the *largest* integer such that $\lfloor \frac{8k+3}{6} \rfloor \leq \lambda(A)$ and $\delta \in \{0, 1\}$.

For $\delta = 0$, Theorem 2 ensures the existence of k edge disjoint A -Steiner trees. Thus,

$$\lambda(A) = \left\lfloor \frac{8k+3}{6} \right\rfloor \leq \frac{8k+3}{6}$$

Which leads to

$$k \geq \frac{6\lambda(A) - 3}{8} \geq \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor \quad (2)$$

For $\delta = 1$, Theorem 2 also ensures the existence of at least k edge disjoint A -Steiner trees, where $\lambda(A) = \lfloor \frac{8k+3}{6} \rfloor + 1$, from which we first obtain

$$k < \frac{6\lambda(A) - 3}{8} \quad (3)$$

we also obtain the following lower bound on k

$$k \geq \frac{6(\lambda(A) - 1) - 3}{8} \quad (4)$$

Upon combining (3) and (4), k can be bounded as,

$$\frac{6\lambda(A) - 3}{8} - \frac{6}{8} \leq k < \frac{6\lambda(A) - 3}{8}$$

Noting that

$$\frac{6\lambda(A) - 3}{8} = \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor + \frac{\Delta}{8}, \quad \Delta \in S \subset \{1, 2, \dots, 7\} \quad (5)$$

(More specifically, it can be shown that for any integer $\lambda(A)$, i.e. δ can be 0 or 1, $\Delta \in \{1, 3, 5, 7\}$, see part-A of the Appendix). Thus the inequality for k becomes

$$\left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor - \frac{6 - \Delta}{8} \leq k < \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor + \frac{\Delta}{8} \quad (6)$$

Note that k is an integer, hence the previous inequality indicates that $k = \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor$ except for $\Delta = 7$, where the inequality has no valid solution. Therefore, our task is to prove that $\delta = 1$ dismisses $\Delta = 7$. To prove this, note that for $\delta = 1$ we have $\lambda(A) = 4q + 2$, $q \in \{0\} \cup \mathbb{Z}^+$. Hence,

$$\begin{aligned} \frac{6\lambda(A) - 3}{8} &= \frac{3 \times 8q + 9}{8} \\ &= 3q + 1 + \frac{1}{8} \implies \Delta = 1 \end{aligned}$$

Which shows that Δ is always 1 for $\delta = 1$. From this we conclude that the solution presented earlier is a valid one. Combining this result with the one obtained for the case of $\delta = 0$ in (2), we can write

$$k \geq \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor, \quad \text{for any } \lambda(A) \in \mathbb{Z}^+ \quad (7)$$

Sending one symbol on each tree, an integral routing throughput of k is achievable. Since π_i is the supreme of all achievable rates, thus $\pi_i \geq k$ and the first bound results. If half-integer routing is allowed, we can up scale the edge capacity by 2 and then divide the integer part of the result by 2. Since this is an achievable routing rate, it represents a lower bound on the half integer routing capacity and thus proves the second part of the theorem. For fractional routing, we upscale edge capacities by n and then divide the integer part of the result by n , i.e. $\pi_f \geq \frac{3\lambda(A)}{4} - \frac{3+\Delta}{8n}$. For sufficiently large n , the second part approaches zero and the third bound results. ■

Corollary 1: For a three terminal multicast network, the coding gain is bounded as

$$\begin{aligned} \mathfrak{G}_i &\leq \frac{\lambda(A)}{\left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor} \\ \mathfrak{G}_{\frac{1}{2}} &\leq \frac{2\lambda(A)}{\left\lfloor \frac{12\lambda(A) - 3}{8} \right\rfloor} \\ \mathfrak{G}_f &\leq \frac{4}{3} + \epsilon' \end{aligned}$$

Proof: $\mathfrak{G} = \frac{\gamma}{\pi} \leq \frac{\lambda(A)}{\pi} \leq \frac{\lambda(A)}{\text{Lower bound on } \pi}$. Applying the appropriate lower bound from theorem 1 and the corollary results.

From corollary 1, it can be seen that the coding gain is always bounded. More specifically, \mathfrak{G}_i is bounded by 3 for

$\lambda(A) \geq 2$ while $\mathfrak{G}_{\frac{1}{2}}$ is strictly less than 2 for $\lambda(A) > 2$. The bound for \mathfrak{G}_f indicates that with fractional routing, 75% of the throughput achievable by coding is always achievable via routing.

B. Multicast Networks with Arbitrary Number of Terminals

Let G be a multicast network modeled as an undirected graph $G(V, E)$ with a set of $N - 1$ receivers $\mathcal{R} = \{R_1, \dots, R_{N-1}\}$ and thus a set of terminals $A = s \cup \mathcal{R} \subseteq V$. The set $X = V - A$ represents the set of non-terminal (relay) nodes in G .

Lemma 1: If G^x , the graph obtained from G by performing a suitable complete splitting at $x \in V - A$, contains a system of k edge disjoint A -Steiner trees, then G contains k edge disjoint A -Steiner trees.

Proof: Let \mathcal{T} be a set of k disjoint A -Steiner trees in G^x . If $d(x) = 2$, let w be the added edge after splitting off edges e and f . For any $T \in \mathcal{T}$, if T contains w then $T - w$, the edges e, f and the vertex x form an A -Steiner tree T' in G and the set $(\mathcal{T} - T) \cup T'$ forms a set of k edge disjoint A -Steiner trees in G . If no $T \in \mathcal{T}$ contains w , then \mathcal{T} is a set of A -Steiner trees in G . If $d(x) \neq 2$, let W be the set of splitting edges and $e(w), f(w)$ be the splitted pair by w . Also let $\mathcal{T}_W = \{T \in \mathcal{T} : T \cap W \neq \emptyset\}$. For every $T \in \mathcal{T}_W$, the tree T' formed by $T - W$, $\bigcup_{w \in W \cap E(T)} g_w$ and x is an A -Steiner tree in G , where $g_w \in \{\emptyset, \{e(w)\}, \{f(w)\}, \{e(w), f(w)\}\}$. Let \mathcal{T}' be the set of such trees, then $(\mathcal{T} - \mathcal{T}_W) \cup \mathcal{T}'$ forms a set of k edge disjoint A -Steiner trees in G .

Alternative proof: Since splitting off does not increase connectivity and complete splitting is a series of splitting off's then $\lambda_{G^x}(u, v) \leq \lambda_G(u, v)$ for all distinct $u, v \in V(G) - x$, with equality if and only if the splitting is suitable. Thus, if G^x contains k -edge disjoint A -Steiner trees, then G does as well.

The following theorem provides a lower bound on the fractional routing capacity, π_f .

Theorem 3: For a multicast network represented by an undirected graph $G(V, E)$ with a set of terminals $A \subseteq V$

$$\pi_f \geq \frac{\lambda(A)}{2} \left(\frac{|A|}{|A| - 1} \right) - \epsilon$$

where ϵ is arbitrarily small for arbitrarily large n and $\lambda(A)$ is the connectivity of A in G .

Proof: The proof uses the following two theorems

Theorem 4 (Kriesell, [10]): For any $l, k \geq 2$, let $G(V, E)$ be an $\lfloor 2k\frac{l-1}{l} + \frac{l-2}{l} \rfloor$ -edge connected graph on l vertices. Then there exist a system of k edge disjoint spanning trees in G .

Theorem 5 (Frank, [11]): Let $G(V + x, E)$ be a connected graph on $V + x$, $d(x) \neq 3$ and no cut-edge is incident to x . Then there exist $\lfloor \frac{d(x)}{2} \rfloor$ disjoint admissible pairs at x .

Consider a complete splitting at every relay node $x \in V - A$. Since G is 2-edge connected, then there are no cut-edges in G . From theorem 5, if $d(x)$ is even then all incident edges to x can be partitioned into disjoint admissible pairs. Since half integer routing is allowed, we can make $d(x)$ even for every $x \in V - A$ by upscaling the capacities of the edges of G

by 2 and then downscaling the solution by 2. Thus a suitable complete splitting exists at all relay nodes in G . Let G' be the graph obtained after such splitting, then $V(G') = A$ and an A -Steiner tree in G' is a spanning tree in G' .

Next we prove that $\lambda(A)$ being any integer, it can be written as

$$\lambda(A) = \left\lfloor \frac{2k(|A| - 1) + |A| - 2}{|A|} \right\rfloor + \delta \quad (8)$$

for any $|A| \geq 2$ and some integer k , where $\delta \in \{0, 1\}$. Let $f_{|A|}(k) = \left\lfloor \frac{2k(|A| - 1) + |A| - 2}{|A|} \right\rfloor$, it is easy to check that $f_{|A|}(k+1) - f_{|A|}(k) = 2 + \left\lfloor \frac{|A| - 2}{|A|} \right\rfloor$, $\Delta \in \{0, 1, \dots, |A| - 1\}$. Noting that $\left\lfloor \frac{|A| - 2}{|A|} \right\rfloor \in \{-1, 0\}$ then, $f_{|A|}(k+1) - f_{|A|}(k) \in \{1, 2\}$. If there exist an integer k such that $\lambda(A) = f_{|A|}(k)$ then $\delta = 0$ and we are done. Otherwise, we choose k as the largest integer such that $\lambda(A) > f_{|A|}(k)$. By the choice of k , $\lambda(A) < f_{|A|}(k+1)$ and since $f_{|A|}(k+1) - f_{|A|}(k) \leq 2$ for any $|A|$ then $\lambda(A) = f_{|A|}(k) + 1$. Thus, for any integer $\lambda(A)$, we can write $\lambda(A) = f_{|A|}(k) + \delta$, $\delta \in \{0, 1\}$, proving the claim in (8).

For $\lambda(A) = \left\lfloor \frac{2k(|A| - 1) + |A| - 2}{|A|} \right\rfloor + \delta$, theorem 4 ensures the existence of k edge disjoint spanning trees in G' , where for $\delta = 0$ we have $k \geq \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor$. If $\delta = 1$, then $\lambda(A) = \left\lfloor \frac{2k(|A| - 1) + |A| - 2}{|A|} \right\rfloor + 1 > \frac{2k(|A| - 1) + |A| - 2}{|A|}$ which results in $k < \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor + \frac{\Delta'}{2(|A| - 1)}$, $\Delta' \in \{0, 1, \dots, 2(|A| - 1) - 1\}$. Also $\lambda(A) - 1 = \left\lfloor \frac{2k(|A| - 1) + |A| - 2}{|A|} \right\rfloor \leq \frac{2k(|A| - 1) + |A| - 2}{|A|}$ results in $k \geq \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor + \frac{\Delta' - |A|}{2(|A| - 1)}$. Upon combining the above two limits of k , we can write $\left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor + \frac{\Delta' - |A|}{2(|A| - 1)} \leq k < \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor + \frac{\Delta'}{2(|A| - 1)}$ which indicates that a valid solution for k exists as long as $\Delta' \in \{1, 2, \dots, |A|\}$. But $\delta = 1$ implies that that this is the case (part-B of the Appendix). Thus for $\delta = 1$, $k = \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor$. Combining this with the result obtained for $\delta = 0$, we conclude that for any integer $\lambda(A) \geq 2$, the number of edge-disjoint A -Steiner trees in G' is $k \geq \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor$.

By Lemma 1, if G' contains k edge disjoint A -Steiner trees then G does, provided that a suitable complete splitting exists at every relay node $x \in V - A$. As it was shown at the beginning of this argument, such a splitting exist if we upscale the capacity of every edge in G by 2 and then downscale our result by 2. Thus, we conclude that G contains k edge disjoint A -Steiner trees, where $k \geq \frac{1}{2} \left\lfloor \frac{2|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor$. Since this represents an achievable throughput, it also represents a lower bound on the half integer routing capacity. For fractional routing, we upscale the edge capacities by n and divide the result by n and for sufficiently large n we obtain $\pi_f \geq \frac{\lambda(A)}{2} \left(\frac{|A|}{|A| - 1} \right) - \epsilon$. ■

Corollary 2: For a multicast network represented by an undirected graph $G(V, E)$ with a set of terminals $A \subseteq V$

$$\mathfrak{G}_f \leq 2 \left(\frac{|A| - 1}{|A|} \right) + \epsilon'$$

Proof: Proof is similar to the one of corollary 1.

C. Examples

Example 1: For the multicast network in Fig.1(a), $\lambda(A) = 2$. The graph resulting from upscaling edge capacities by 2 and then performing suitable complete splitting at all relay (non-hatched) nodes is shown in Fig.1(b), where thick edges have double the capacity of normal ones. G' is a broadcast network, and from theorem 3 with $A = V(G')$ the routing capacity of G' can be bounded as $\pi_f \geq \frac{4}{3}$. From theorem 3, this also serves as a lower bound on the routing capacity of G and thus the coding gain in G is bounded by $\mathfrak{G}_f \leq 1.5$. Note that $\frac{4}{3}$ is a lower bound on the capacity, a routing scheme achieving a throughput of 1.5 for G' using fractional routing with $h = 9$ and $n = 6$ is shown in Fig.1(b). For the network on G , a fractional routing throughput of 1.8 is possible with $h = 9$ and $n = 5$, Fig.1(a).




Fig. 1. Example 1: (a)-Multicast Network G . Also shown a fractional routing scheme that achieves a throughput of 1.8. The source has $h = 9$ messages $\{a_0, \dots, a_8\}$ and the capacity of each link is upscaped by $n = 5$. The symbols on the edges represent the messages carried by the edges and the arrows indicate the direction of the flow. All sinks will recover all the messages and thus a throughput of $\frac{9}{5} = 1.8$ is possible. (b) Split graph G' together with a fractional routing scheme that achieves a throughput of $\frac{9}{6} = 1.5$.

Example 2: Fig.2(a) represents a multicast network with a set of terminals $A = \{v_0, v_1, \dots, v_{|A|-1}\}$ and a set of relay nodes $X = \{x_1, \dots, x_{|X|}\}$ where $X = V - A$ and v_0 is assumed to be the source node. The graph G' resulting from performing suitable complete splitting at all relay nodes is shown in Fig.2(b). From theorem 3 with $\lambda(A) = 2$, the routing capacity can be lower bounded as $\pi_f \geq \left(\frac{|A|}{|A| - 1} \right)$ for both G and G' .

Next we show that for the network G in Fig.2(a), $\gamma = \pi_f = \left(\frac{|A|}{|A| - 1} \right)$. From [1], $\gamma \leq \eta_G(A)$ where $\eta_G(A)$ is the edge strength defined as

$$\eta_G(A) = \min_{\mathcal{P}} \frac{E_G(\mathcal{P})}{|\mathcal{P}| - 1}$$

and the minimization is over all possible partitions $\mathcal{P} = \{V_0, V_1, \dots, V_{|\mathcal{P}|-1}\}$ of $V(G)$ such that each component, V_i , of the partition \mathcal{P} contains at least one terminal, i.e. $V_i \cap A \neq \emptyset$. $E_G(\mathcal{P})$ is the total capacity of edges between distinct components. Let us choose a partition \mathcal{P} such that $V_i \cap A = v_i \forall i \in \{0, 1, \dots, |A| - 1\}$. For such a partition,

$|\mathcal{P}| = |A|$. Also since G is a cycle with each edge having unit capacity, then $E_G(\mathcal{P}) = |A|$. Because $\eta_G(A)$ is the minimum of the ratio $\frac{E_G(\mathcal{P})}{|\mathcal{P}|-1}$ over all possible partitions \mathcal{P} , then $\gamma \leq \eta_G(A) \leq \frac{|A|}{|A|-1}$. Combining this with the lower bound on π_f obtained earlier from theorem 3 and the fact that $\pi_f \leq \gamma$, we obtain

$$\frac{|A|}{|A|-1} \leq \pi_f \leq \gamma \leq \frac{|A|}{|A|-1}$$

Thus, for the multicast network G shown in Fig.2(a), $\pi_f = \gamma = \frac{|A|}{|A|-1}$.

A routing scheme achieving the fractional capacity can be advised as follows: Let $h = |A|$ and $n = |A| - 1$. Node v_i forwards the $|A| - (i + 1)$ symbols $a_0, a_1, \dots, a_{|A|-1-(i+1)}$ to terminal v_{i+1} , $\forall i \in \{0, 1, \dots, |A| - 2\}$. Terminal v_0 sends $|A| - 1$ symbols, $a_1, \dots, a_{|A|-1}$, to $v_{|A|-1}$. Finally, terminal v_{i+1} forwards i symbols, $a_{|A|-i}, \dots, a_{|A|-1}$, toward terminal v_i , $\forall i \in \{1, \dots, |A| - 2\}$. On the other hand, all relay nodes in X do nothing but forwarding whatever they receive at their input edge to their output edge, Fig.2(c). Note that with such a routing scheme, each edge carries $|A| - 1$ symbols and each terminal receives $|A|$ symbols. Achieving a routing throughput of $|A|/(|A| - 1)$. Fig.2(c), shows an example with 5 terminals, $\{v_0, v_1, v_2, v_3, v_4\}$, and 2 relay nodes, $\{x_1, x_2\}$.

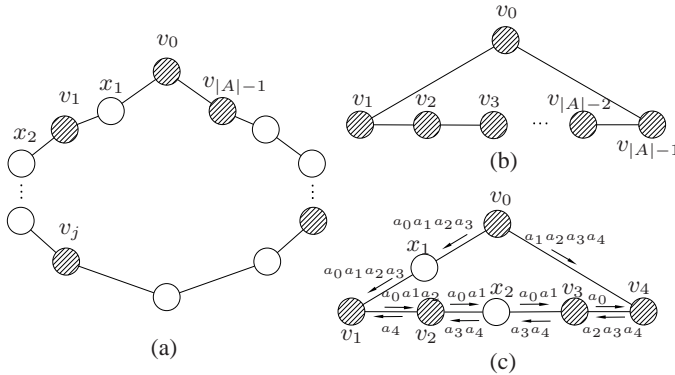


Fig. 2. Example 2: (a)-Multicast Network G , (b) Splitted graph, G' , (c) An instance of G with 5 terminals and 2 relay nodes.

APPENDIX

We use the notation $(a)_b$ to denote a modulo b

A. $\Delta \in \{1, 3, 5, 7\}$ for the multicast network with three terminals.

Proof: From (1) we can write $\lambda(A) = 4q + c$, $c \in \{0, 1, 2, 3\}$. Also from (5)

$$\frac{6\lambda(A) - 3}{8} = \left\lfloor \frac{6\lambda(A) - 3}{8} \right\rfloor + \frac{\Delta}{8}$$

In other words,

$$\Delta = (6\lambda(A) - 3)_8 = (3 \times 8q + 6c - 3)_8 = (6c - 3)_8$$

thus,

$$\Delta = \begin{cases} 5, & c = 0 \\ 3, & c = 1 \\ 1, & c = 2 \\ 7, & c = 3 \end{cases}$$

Note that for the case $\delta = 1$, then $c = 2$ and thus, $\Delta = 1$ as it was shown before.

B. For the multicast network with any number of terminals, $\Delta' = 0$ if and only if $\Delta = 0$ and $\delta = 0$.

Proof: From (8)

$$\lambda(A) = \frac{2k(|A| - 1) + (|A| - 2)}{|A|} - \frac{\Delta}{|A|} + \delta$$

From which we can write,

$$\begin{aligned} k &= \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} + \frac{\Delta - |A|\delta}{2(|A| - 1)} \\ &= \left\lfloor \frac{|A|\lambda(A) - |A| + 2}{2(|A| - 1)} \right\rfloor + \frac{\Delta' + \Delta - |A|\delta}{2(|A| - 1)} \end{aligned}$$

where $\delta \in \{0, 1\}$, $\Delta \in \{0, 1, \dots, |A| - 1\}$ and $\Delta' \in \{0, 1, \dots, 2(|A| - 1) - 1\}$. This shows that $\Delta' + \Delta - |A|\delta$ is divisible by $2(|A| - 1)$. Thus,

$$(\Delta' + \Delta)_{2(|A|-1)} = |A|\delta$$

From this we note that if $\Delta = 0$ and $\delta = 0$, then $\Delta' = 0$, which proves the 'if' part of the claim. Conversely, if $\Delta' = 0$, then $\Delta = |A|\delta$, which shows that if $\delta = 1$, then $\Delta = |A|$, a contradiction (since $\Delta \in \{0, 1, \dots, |A| - 1\}$). Thus, $\delta = \Delta = 0$ and the claim follows.

REFERENCES

- [1] Z. Li and B. Li, "Network coding in undirected networks," in *Proc. 38th Annu. Conf. Information Science and Systems (CISS)*, Princeton, NJ, Mar. 2004, pp. 257–262.
- [2] R. Ahlswede, N. Cai, S.-Y. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. on Info. Theory*, vol. 46, no. 4, pp. 1204–1216, July 2000.
- [3] S. Jaggi, P. Sanders, P. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Trans. on Info. Theory*, vol. 51, no. 6, pp. 1973 – 1982, June 2005.
- [4] C. Chekuri, C. Fragouli, and E. Soljanin., "On average throughput and alphabet size in network coding," *IEEE Trans. on Info. Theory*, vol. 52, No. 6, pp. 2410–2424, June 2006.
- [5] A. Agarwal and M. Charikar, "On the advantage of network coding for improving network throughput," in *Proc. 2004 IEEE Information Theory Workshop*, San Antonio, Texas, Oct. 2004, pp. 247–249.
- [6] L. Zosin and K. Khuller, "On directed Steiner trees," in *Proc. 13th Annu. ACM/SIAM Symp. Discrete Algorithms (SODA)*, San Francisco, CA, January 2002, pp. 59–63.
- [7] E. Halperin, G. Kortsarz, R. Krauthgamer, A. Srinivasan, and N. Wang, "Integrality ratio for group Steiner trees and directed Steiner trees," in *Proc. 14th Annu. ACM/SIAM Symp. Discrete Algorithms (SODA)*, Baltimore, MD, Jan. 2003, pp. 275–284.
- [8] S. Riis, "Linear versus non-linear Boolean functions in network flow," Tech. Rep., 2003. [Online]. Available: <http://nick.dcs.qmul.ac.uk/~smriis>.
- [9] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear network coding," *IEEE Trans. on Info. Theory*, vol. 49, no. 2, pp. 371–381, Feb. 2003.
- [10] M. Kriesell, "Edge-disjoint trees containing some given vertices in a graph," *Journal of Combinatorial Theory*, ser. B 88, pp. 53–65, 2003.
- [11] A. Frank, "On a theorem of mader," *Discrete Mathematics*, 101, pp. 49–57, 1992.